Powering the Data Center - Distributed 48Vdc is an Effective Choice

Dan McMenamin
President – Dan McMenamin and Associates, Inc.

INTRODUCTION

Effective management takes creativity. The engineer, as both manager and steward of his company's resources is called upon to deliver world-class service and reliability, usually on a compressed timeline, funded from a parsimonious budget such that every Dollar, Euro or Pound is wrung from the fingers of a soulless bean-counter screeching "first cost" from his or her iron cubicle. Accomplishing the engineer's job in the pressure cooker world of Information Technology (IT.) takes a creative person with the guts to think, design and manage outside the box. While the centralized, decentralized architecture pendulum has swung back and forth in telecom for quite awhile now, the decentralized concept is somewhat foreign to data centers. Similarly, while 48V data centers have gained traction, the concept of distributed 48V power to data center systems hasn't seen a lot of press but offers significant advantages over centralized systems.

One way to look at distributed 48Vdc power is a microgrid or series of microgrids bringing modularity and cost containment to the task of supplying energy to the Routers, Servers and other inner-workings that comprise the I.T. world.

This paper will explore the many conditions and parameters where distributed low voltage dc makes solid engineering sense in the I.T. environment.

Failure Impact

As the demand decreases for wireline telephony and dramatically increases for wireless, the need for power availability is greater than ever. Wireline central offices are limited to a given territorial radius by the loop length of subscriber cable pairs. In the digital world of wireless telephony, that restriction went away. Therefore, where a wireline central office might have served 50,000 or even 90,000 telephone customers, a wireless Mobile Switching Center (MSC) easily can handle the traffic for a half-million lines or more. Therefore, the impact of power problems looms large in the industry.

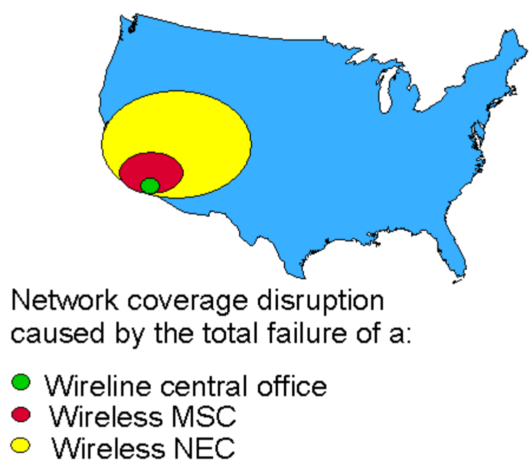

Figure 1. a comparative view of the impact of power failure in telephone facilities

With the advent of highly specialized wireless data services, carriers had to increase the bandwidth of their network between data providers and cell towers and from the cell site to the customer handset. The traditional T-1 line was far too limited in capacity and thus, the Internet protocol (I.P.) Ethernet based systems came into prominence. Much wireless telephone traffic now passes through what essentially are central offices filled with data center type systems. Some operating companies call such facilities, Network Equipment Centers (NEC) in order to differentiate them from traditional data centers because very different fire protection standards apply to them. Traditional data centers are governed by NFPA 75[2], whereas telecommunications facilities greater than 1,500 square feet (139.4 square meters) are governed by NFPA 76[3].

AC UPS

Protected power for most data centers is provided by a centralized ac Uninterruptible Power Supply (UPS) as shown in Figure 2, or a parallel redundant ac UPS arrangement such as is shown in Figure 3.

One advantage of using transformer equipped Power Distribution Units (PDU) is that for both 208 and 480 volt UPS systems, any common mode electrical noise induced into the cabling between the UPS and the PDU is consumed in the delta primary winding of the transformer where its energy simply becomes a trivial heat source. For 480 volt UPS systems there is the added advantage of using smaller conductors to the PDU and then stepping down to 208/120 to feed single and three phase loads.

The redundant arrangement in Figure 2 utilizes what's called either a 'paralleling cabinet" or a 'system control cabinet' to couple the output from two or more UPS modules to Power Distributing Units (PDU)'s and provide both a static switched bypass source or a manually switched maintenance bypass.

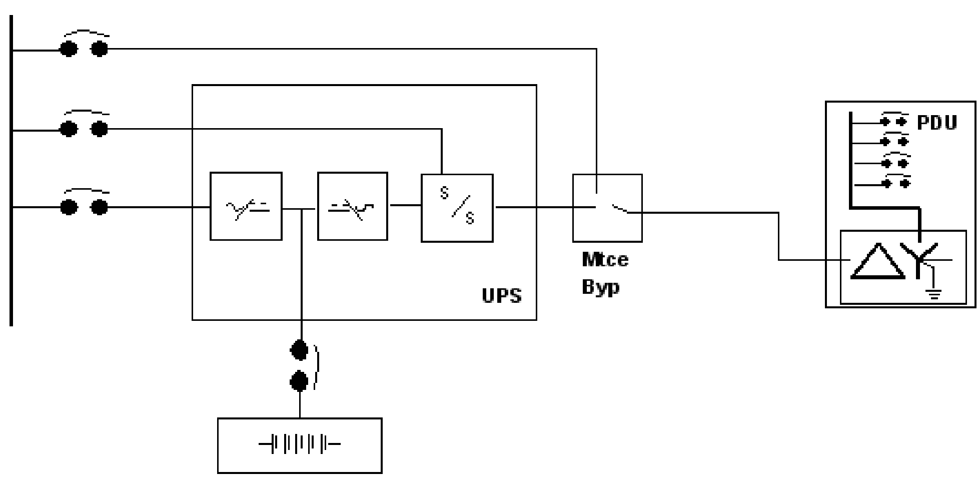

Figure 2. Block diagram of a traditional double-conversion non redundant ac UPS system

These configurations have been around since the earliest days of the data center business and they comprise the power scheme for most data centers around the world although the actual operating voltages vary from country to country.

In the 1990’s as Information Technology (IT) equipment began making inroads into the dc powered world of telecommunications central offices, a bit of stress and controversy came with them. At the time, virtually all IT equipment was powered by ac and the telecom network was powered by a rather bulletproof 48 Volt dc rail. While the applications and features to be gained with the addition of IT network elements were a superb value, powering the units was less so. The telecom industry has long used inverters for ac powered items, but doing so adds cost; and to some extent reduces reliability by adding system complexity and points of failure. Some telephone companies began pushing their IT suppliers for 48 Volt powered versions of their products and before long, companies like Cisco and others responded accordingly.

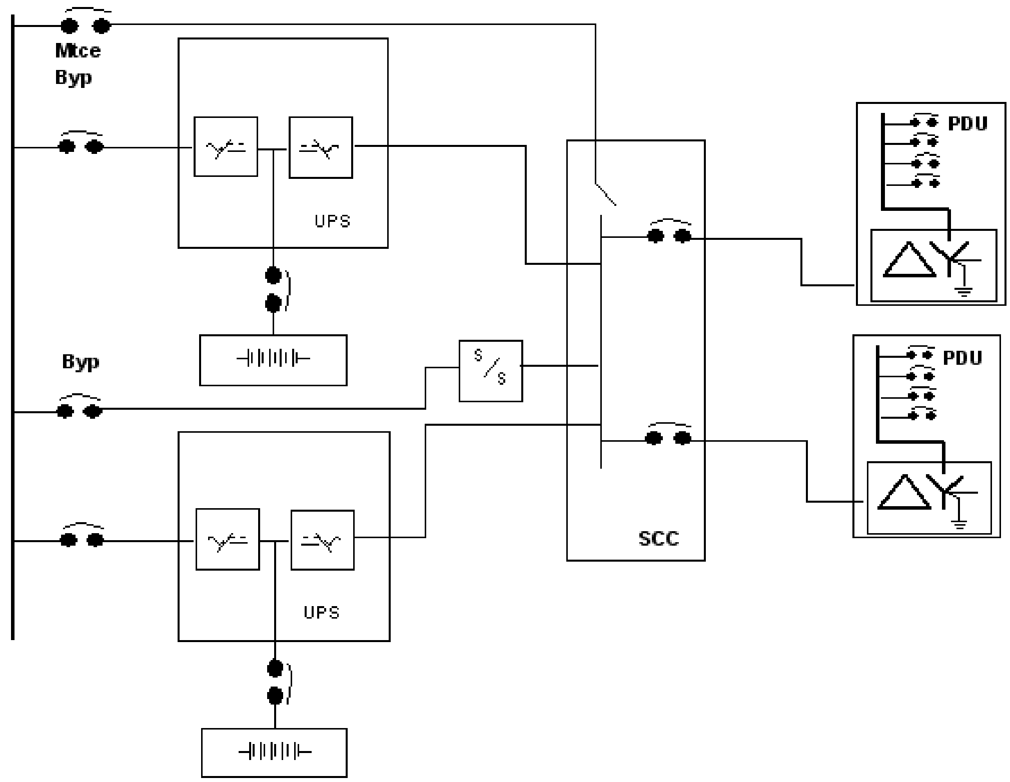

Figure 3. Block diagram of a traditional double-conversion redundant ac UPS system

Sometime in 1998, a meeting was held in Stockholm among the Technical Subgroup on Telecommunications Energy Systems of the Power Electronics Society of the Institute for Electrical and Electronics Engineers. The product of that meeting was a White Paper titled “-48VDC Computer Equipment Topology – An Emerging Technology.” [1]

The intent of the paper was a call to the industry to provide IT solutions for use in the telecommunications network that would efficiently continue operation during an ac mains failure of four (4) to eight (8) hours and that the 48 volt bus was the most effective power source for such an operation. The paper cited a number of case studies as well as providing a model that showed a traditional telecommunications 48 volt plant as having some twenty times improved availability over that of ac UPS. In essence, the paper reignited the hard fought “War of the Currents,” waged in the late 1800’s between Thomas Edison and his former employee Nikola Tesla, who, at the time worked for George Westinghouse.

As IT systems and subsystems made greater penetration into telecommunications network infrastructure, some operating companies began to plan entire data centers powered by 48 Volts with a handful of marginal ac powered systems fed from centralized power room inverters or point of use inverter units. The advantages of a 48 Volt data center were many including the facts that one didn't have to balance phase load across the Wye (Secondary) windings of PDU transformers or use electricians to add minor circuits. One disadvantage is that the deployment of large scale, energy-dense IT systems produced heat loads that drove companies to use access floors as a cooling air plenum. Plenum floors require plenum-rated cables, and while various companies manufacture communications and data cabling for such applications, no one makes a dc power cable suitable for deployment in an air plenum. Further, the cabling volume was so large that it wouldn't easily fit into traditional conduits and wiring troughs would be so large as to obstruct the flow of fire suppression gasses used in 'clean-agent' systems. Another disadvantage is the commodity and installation cost of large copper cables needed to overcome voltage drop between a centralized power room and the systems to be operated.

Distributed 48 volt deployments have seen limited use, however, the size, weight and occasional leakage problems associated with lead acid batteries has served to deter most telco engineers from such an arrangement. Such is true even of Valve Regulated Lead Acid (VRLA) cells even though they have about a 20% improved power density than Vented (formerly called Flooded) cells. The improved density comes from two factors: a. there are thinner plates resulting in more power at reduced life; and b. less electrolyte is used in the cell construction. Emerging battery technologies promise densities significantly higher than their lead acid counterparts. Lithium-ion and sodium-nickel cells may have the potential to make a distributed architecture more efficacious for the central office or data center environment simply because they occupy less floor space.

There are few, if any, data centers that operate as a single complex of servers, routers and related equipment. Generally, there are communities of systems and subsystems sharing a conditioned space with a centralized or decentralized source of protected power.

Figure 4. Rows of UPS fed Power Distributing Units (gray cabinets) feeding power to rows of IT equipment (Black cabinets)

Wireless telecom facilities, for example, house groups of systems serving voicemail storage, small message (text) processing and storage and various other systems such as home location registers, message registers etc., that interconnect with each other but there is no particular requirement for a single power source. If a voice mail system was depowered that would be an inconvenience but not drastically curtail telecommunications services. Indeed, many central offices, mobile service centers and the like, designed under fire protection standards prescribed by NFPA 76, are arranged for selective zone depowering. Using this protection scheme, fire departments are able to remove power to certain equipment areas without depowering other areas of the building. The intent of such an arrangement is to minimize the impact of a depowering on the vital telecommunications network. Such an arrangement is quite different from the 'kill everything', Emergency Power Off (EPO) used in an IT world.

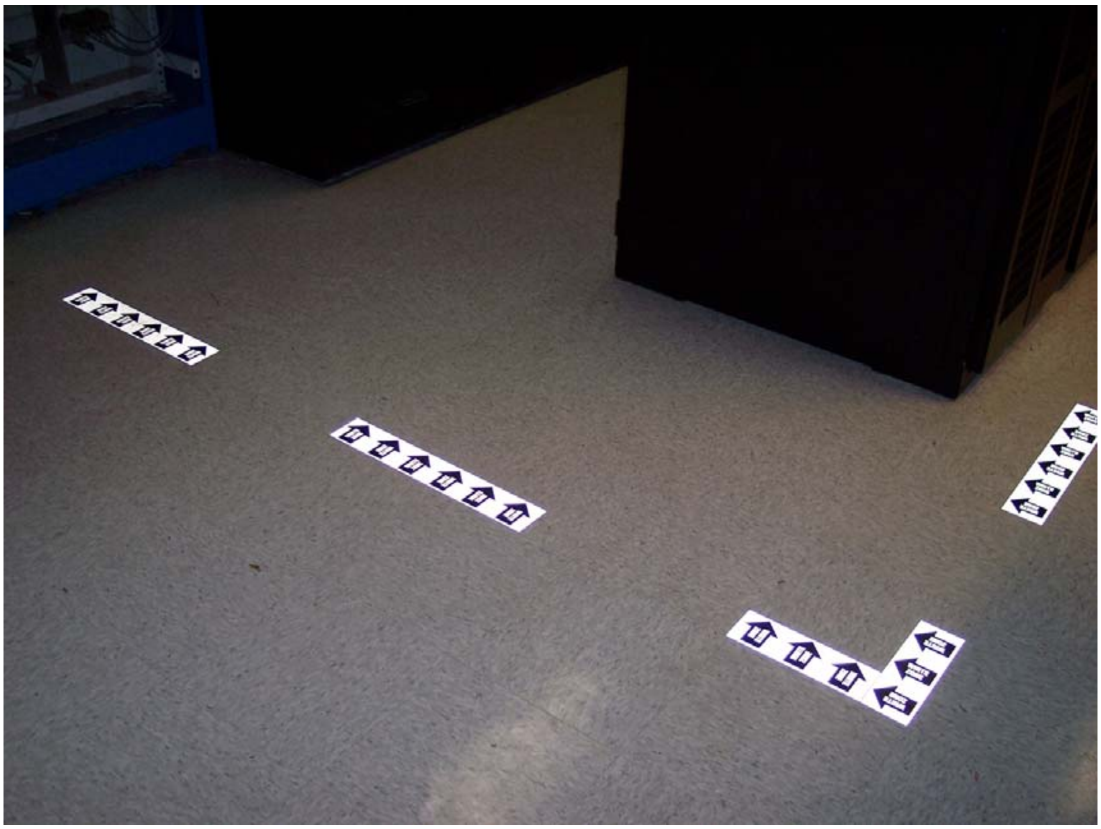

Figure 5. Equipment aisles identified with color coded reflective decals for zoned emergency power disconnect

Because most server equipment has dual A/B power supplies sized for N+N redundancy at the shelf level or individual server, the architecture lends itself to pairs of distributed -48 volt dc plants. Following the zoned disconnect scheme that has found favor with fire departments across the nation, equipment communities might be arranged similar to that shown in Figure 6 with the A dc plant feeding the A sides of all equipment and the B dc plant feeding all the B sides.

With this level of redundancy, yet another fire service tool and equipment reliability initiative could be built into the network relatively inexpensively. Under zoned disconnect plans, fire service responders locate a bay that is emitting smoke and they note what color floor decals surround the problematic bay or cabinet. Then, they follow color coded floor decals to the dc power room and then to the distribution fuse or circuit breaker bay(s) feeding that color zone. Then, the fire service simply pulls those fuses or trips those circuit breakers to kill power to that zone. With all other equipment zones intact, power-wise, the impact on telecommunications services may be minimized. Quite possibly, the impairment will be severe but at least minimized as opposed to a total power-down.

If A/B pairs of distributed architecture plants fed individual equipment zones, then buttons protected against accidental operation could be placed at the equipment room door to facilitate depowering that zone in a manner that offers high reliability to the network by giving the option to shut down only the A plant or the B plant in a given zone, and both only if need be.

For example, a fire service responds to smoke alarms in an equipment room, a Class C (electrical) fire, and they detect smoke coming from a cabinet in the blue equipment zone. They then would visit a nest of shutdown buttons near the door and operate the shutdown for blue zone A plant. At this point they'd wait a minute or so and see if the smoke had significantly diminished. If the smoke has diminished no further disconnection is necessary. Servers would remain in service fed by the remaining power plant. If the smoke did not diminish, then the fire department would operate the blue zone B plant switch, fully depowering that zone. If the fire was in the dc power bay itself and power off had not resolved the problem, then the traditional floor decals would show the path to kill all ac power and generator power in the facility, placing the remaining equipment groups on battery but still in service.

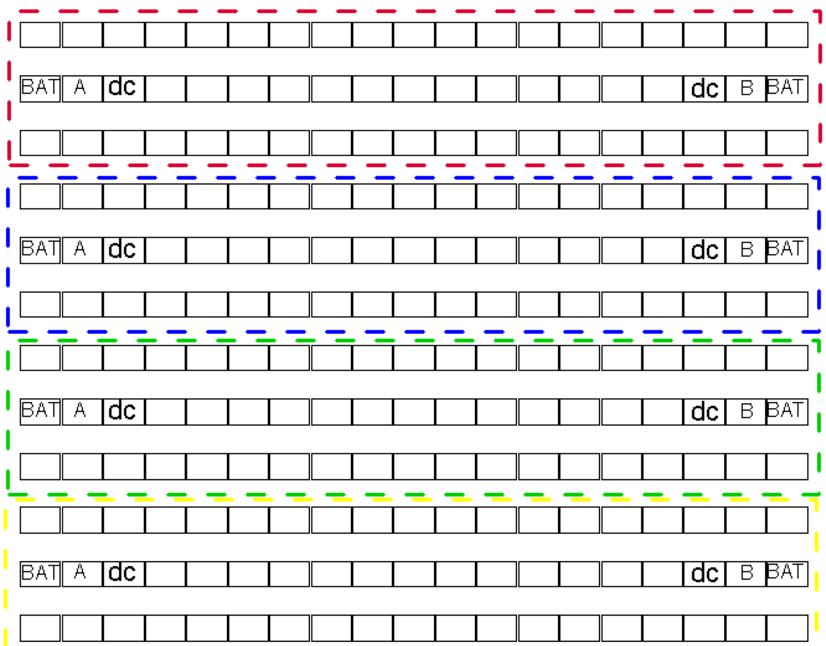

Figure 6. Equipment aisles identified with color coded reflective decals for zoned emergency power disconnect

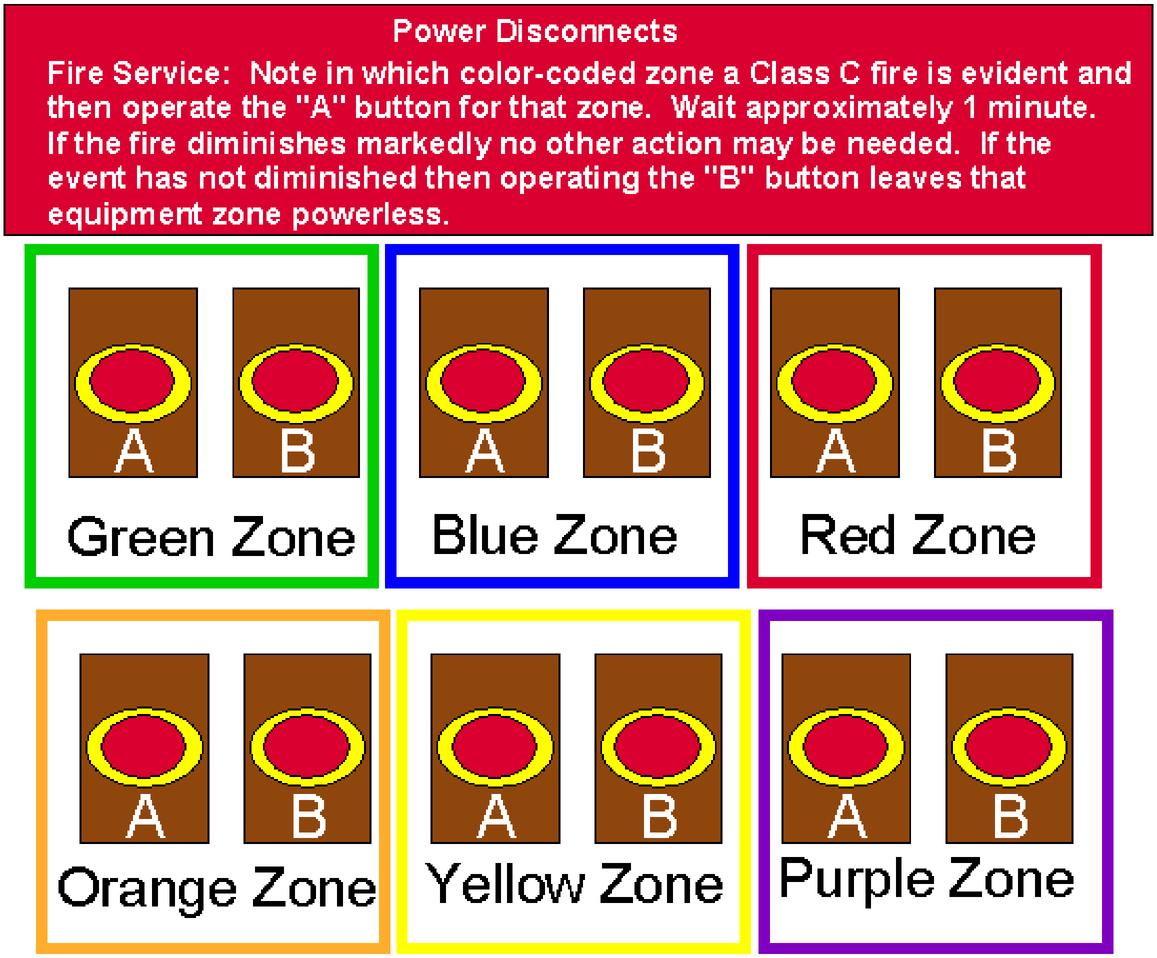

Figure 7. A zoned power disconnect arrangement for facilities with distributed dc (6 zones are shown in this example).

What would the power off circuit look like at the distributed dc bay? Output contactors are not new, nor are shunt trips for main feeder circuit breakers as might feed rectifiers. One issue is that many codes require that power off circuits be tested periodically. If all dc equipment was A/B fed and properly administered, depowering either dc plant should be a non-issue. Of course, errors in A/B assignment where both sides of a server or other subsystem were fed by either plant would evidence themselves pretty quickly as those equipments would fail.

Another approach might be to include a bypass switch around the power off contactor similar to what's shown in Figure 7. Because distributed dc plants would be fairly small, incorporating a molded case bypass switch with auxiliary contacts wouldn't be very expensive. The auxiliary contacts would operate an indicator Light Emitting Diode (LED) and perhaps a piezoelectric buzzer to alert that the plant disconnect was in a bypass mode.

The test plan would be to close the bypass switch for the plant under test and note the LED and annunciator operating. Next, operate the remote power off button for that plant and note that the contactor had opened and the rectifiers depowered. Upon restoring the power off button, the contactor should reengage and at that point the bypass switch should be returned to the open position and the ac power main breaker reset. The rectifiers would restart and the test therefore is completed to satisfy the code without ever dropping the load.

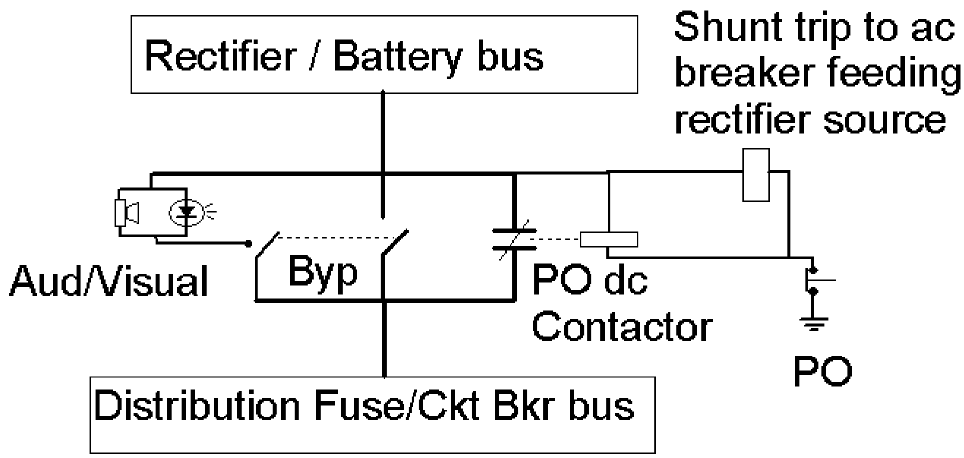

Figure 8. A zoned power disconnect arrangement with a bypass feature for system tests

Because a distributed dc architecture would be close to the equipment it serves the wiring becomes small enough to fit within standard flexible steel conduits eliminating the problem caused by the industry unavailability of plenum rated dc power cable.

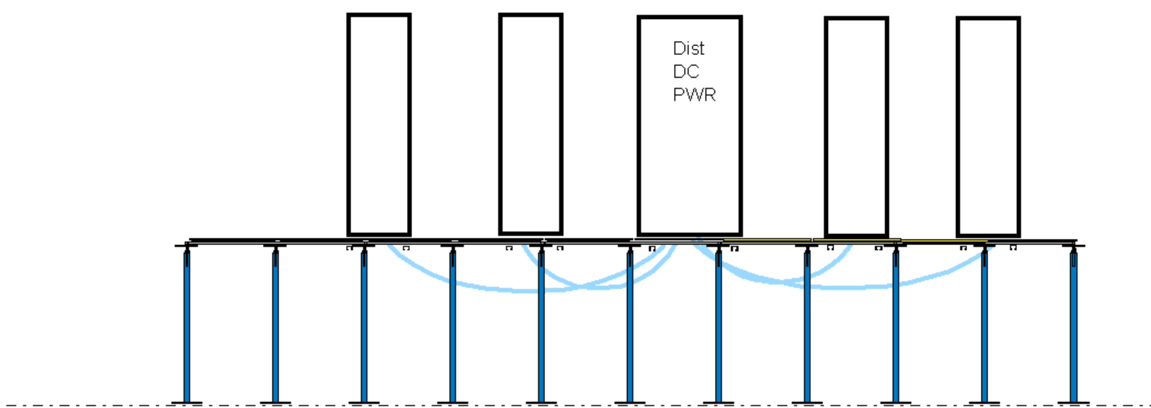

Figure 9. End view sketch of rows showing -48 volt feeding cabinets via flexible conduits.

The Battery

The engineering choice of a battery plays a critical role in the design of any dc system whether centralized or distributed. The battery, for all practical purposes is like a parachute for the load; and is what saves your bacon. The more reliable and the more capacity the battery (or the parachute) has, the slower and gentler you hit the ground.

When designing for an A/B fed dual-plant arrangement the battery reserve calculation becomes complicated because battery reserve time is not a linear load function. As will be explained further, doubling the load does not simply halve the battery reserve time. Also, there are a number of A/B arrangements in telecom and data systems and therefore the design engineer must understand in order to know what will be the load profile with which he or she is working.

Referring to Figure 10, some A/D systems utilize two 48 volt inputs which pass through two 48/5 volt converters which in turn flow through forward biased diodes so that each feed and dc/dc converter carries exactly half the load of the circuit. In the event that either feed or dc/dc converter fails, the remaining feeder 'sees' a load doubling. The forward-biased diodes often are referred to as "ORing" diodes because both "or" one input can carry the load. The diode serves to isolate the inputs so that a fault on the 'A' side doesn't impinge on the 'B' side.

Other A/B arrangements have dual inputs and dual loads, such as processors, and both processors run in parallel. If one input fails that processor fails and the remaining one keeps operating. In this case, the load would not 'double' as each side runs independently.

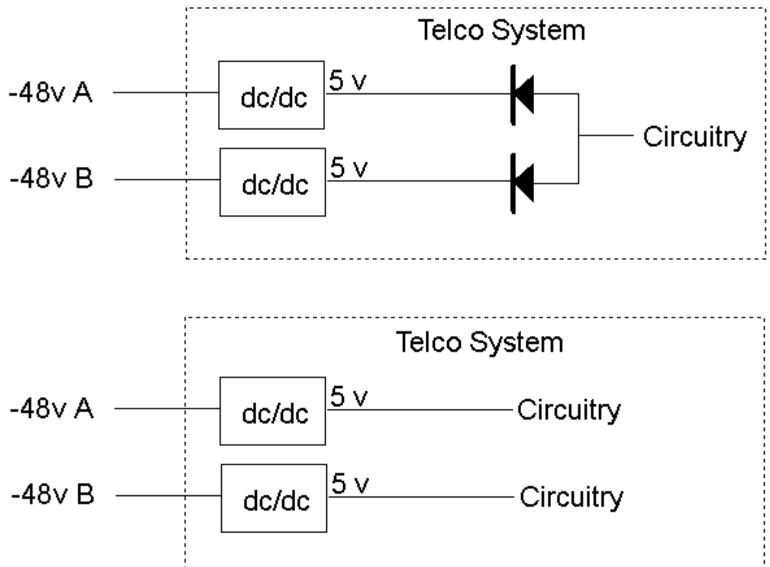

Figure 10. depicts two A/B input systems that have very different load profiles.

These issues are critical to battery design. Figure 11 shows an excerpt from the battery discharge chart for several C&F vented lead acid cells commonly used in telecommunications service. Taking a discharge load of approximately 150 Amperes per string (159 on the chart) on a battery comprised of 1680 Ampere Hour cells, the battery would discharge down to a telco typical minimum volts per cell in 10 hours.

If that load was to double to 300 amperes as would be the case in diode-ORing A/B feeds, the battery reserve time would be 4 hours. It's a common error for people to assume that doubling the load on a 10 hour battery would provide a 5 hour reserve time but it's clear from the discharge chart that such is not the case. This is not a battery vendor-specific issue, but one common to batteries in general.

Final volts	Battery models	Nominal rates at 77F (25C) (includes connector voltage drop) Amperes per cell										
		24 Hr	10 Hr	8 Hr	6 Hr	5 Hr	4 Hr	3 Hr	2 Hr	1.5 Hr	1 Hr	.5 Hr
1.84	2LCT/LT- 840	39	80	94	115	130	150	177	220	250	288	336
	LCT/LT- 840	39	80	94	115	130	150	177	220	250	288	336
	LCT/LT-1008	47	96	113	138	156	180	213	263	299	346	403
	LCT/LT-1176	55	112	132	161	182	209	248	307	349	403	470
	LCT/LT-1344	63	127	150	184	208	239	284	351	399	461	537
	LCT/LT-1680	79	159	183	230	260	299	355	439	499	576	671
	LCT/LT-1848	86	175	207	253	286	329	390	483	549	633	738
	LCT/LT-2016	94	191	225	276	311	359	426	527	599	691	806
	LCT/LT-2176	93	199	233	299	342	402	487	616	706	816	945
	LCT/LT-2320	99	212	251	319	365	429	519	657	753	870	1,008

Figure 11. a battery discharge chart showing the impact of load doubling on reserve time

Accordingly, it is critical that the engineer designing the battery determine the load profile so as to understand exactly how much load will be presented to the battery and under what circumstances so that the proper capacity is designed into the system. Otherwise, the design may fall well short of expectations. Returning to the parachute metaphor, such a miscalculation would be like putting a 250 pound person under a canopy designed for a 150 pound person and his or her ride is likely to end badly.

Cost?

Because of sales volumes, generally, smaller dc plants are more cost competitive than large dc plants intended for central offices. And they are cheaper to ship and install. Together with the reduced cost of copper resulting from such an arrangement, such deployments would offer large savings and even larger gains in overall system flexibility and reliability.

The significantly reduced size of sodium-nickel batteries and lithium-ion cells made available through Electric Vehicle (EV) research, would make such deployments sensible for the equipment room. Both of these technologies behave well in the benign environment of an equipment room. Their cost, while significantly higher than lead acid, would be offset by the per square foot cost of building construction or addition.

Conclusions

Is a distributed -48 volt dc system the answer to each and every data center? Probably not, yet there is every reason to believe that for many data centers and in particular, in avoiding building additions to existing data centers, that the 48 volt option offers engineers the opportunity to win the tug of war between cost and reliability. Those who win that game remain competitive.

References

1. -48 V DC Computer Equipment Topology - an Emerging Technology J. Åkerlund, B. Lindemark, D. McMenamin, S. Muroyama, J. Parsons, J. Poulin, C. Riddleberger, K. Yotsumoto, INTELEC ’98, San Francisco, California USA
2. NFPA 75 Standard for the Protection of Information Technology Equipment. National Fire Protection Association, Quincy MA
3. NFPA 76 Standard for the Fire Protection of Information of Telecommunications Facilities. National Fire Protection Association, Quincy MA